arXiv:1605.06377v1 [cs.LG] 20 May 2016

Towards Automation of Knowledge Understanding: An Approach for Probabilistic Generative Classifiers

Dominik Fisch * Christian Gruhl † Edgar Kalkowski †
Bernhard Sick † Seppo J. Ovaska ‡

After data selection, pre-processing, transformation, and feature extraction, knowledge extraction is not the final step in a data mining process. It is then necessary to understand this knowledge in order to apply it efficiently and effectively. Up to now, there is a lack of appropriate techniques that support this significant step. This is partly due to the fact that the assessment of knowledge is often highly subjective, e.g., regarding aspects such as novelty or usefulness. These aspects depend on the specific knowledge and requirements of the data miner. There are, however, a number of aspects that are objective and for which it is possible to provide appropriate measures. In this article we focus on classification problems and use probabilistic generative classifiers based on mixture density models that are quite common in data mining applications. We define objective measures to assess the informativeness, uniqueness, importance, discrimination, representativity, uncertainty, and distinguishability of rules contained in these classifiers numerically. These measures not only support a data miner in evaluating results of a data mining process based on such classifiers. As we will see in illustrative case studies, they may also be used to improve the data mining process itself or to support the later application of the extracted knowledge.

1 Introduction

Data mining (DM) can be seen as a multi-step process as shown in the data mining pyramid (see Fig. 1) which was introduced by Embrechts et al. in [16]. The idea of this pyramid can briefly be summarized as follows: Raw *data* are pre-processed to condense application-specific *information* in attributes or features. Then, *knowledge* is extracted, e.g., by building classification or regression models. By analyzing this knowledge off-line (i.e., after the model is learned from training data) and by using it in a given application (on-line) it is possible to come to a deeper *understanding* of its working principles and to gain some *experience* in using it, respectively. Both will support the efficient and effective application of the knowledge. Finally, this kind of meta-knowledge (knowledge about knowledge) will eventually help to solve similar kinds of application problems. That is, the final step of *wisdom* is reached by transferring the knowledge to other application domains. While the steps from data over information to knowledge are

*BMW Group, 80788 Munich, Germany (e-mail: dominik.fisch@bmw.de)
†University of Kassel, Department of Electrical Engineering and Computer Science, Wilhelmshoeher Allee 73, 34121 Kassel, Germany (e-mail: {cgruhl,kalkowski,bsick}@uni-kassel.de)
‡Aalto University, Department of Electrical Engineering and Automation, Espoo, Finland (e-mail: seppo.ovaska@aalto.fi)

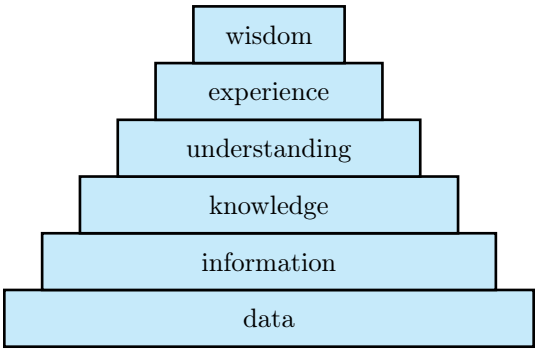

Figure 1: The data mining pyramid (adopted from [16]).

well supported by appropriate algorithms and commercial or free tools, there is still a lack of techniques that support the subsequent steps.

In this article, we address the problem of knowledge understanding by analyzing its properties off-line. We use *classifiers based on probabilistic mixture models* (CMM, see [22, 6]) which can be used in many DM applications. These classifiers can be termed to be hybrid in the sense that different kinds of distributions are combined for different kinds of attributes (e.g., continuous or categorical) which makes the classifier very flexible. CMM are trained from data samples using expectation maximization or related techniques such as variational Bayesian approaches [6]. For CMM, we propose various measures that assess the *informativeness*, *uniqueness*, *importance*, *discrimination*, *representativity*, *uncertainty*, and *distinguishability* of rules contained in this classifier. The measures can be used in different ways: For example, it is possible to prune classifiers, to rank classification rules, or to detect novel kinds of knowledge (cf. anomaly detection) during the classifier's application. However, the proposed measures should be seen as a first yet important step towards an automation of knowledge understanding. Thus, with the assumptions we make, we do not cover all possible aspects of real World applications so far. For instance, other probability distributions are needed for other attribute types.

This article is a substantially extended version of a conference article (cf. [20]). The main contributions compared to the previous article are: All existing measures are revised and we introduce two additional measures (**uncertainty** and **distinguishability**) to further analyze different aspects of generative classifiers. The evaluation of the measures is also completely new and improved. In four case studies utilizing over 20 artificial and real-world benchmark data sets we illustrate how our proposed measures can help data scientists.

In the remainder of the article we first briefly discuss related work in Section 2. Then, we describe the classifier and introduce the various measures in Section 3. Illustrative case studies in Section 4 show how the new measures can be applied. In Section 5 we summarize the key findings and give an outlook to future work.

2 Related Work

Obviously, the measures we are looking for are closely related to so-called *interestingness measures* in data mining. *Data mining (DM)*, today often used as a synonym of *knowledge discovery in databases (KDD)*, deals with the "the nontrivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data" [17]. But, how can this "interestingness" be assessed numerically? Obviously, there are objective facets of interestingness such as validity and subjective facets such as usefulness (see, e.g., [32, 33, 44, 58, 62]).

Objective measures analyze the extracted knowledge without relating it to the users' prior knowledge or needs. These measures are based, e.g., on so-called information criteria or on data-based measurement techniques (see [65] for an overview). Examples for the former are

Akaike's information criterion or the Bayesian information criterion. Examples for the latter are statistical measures such as sensitivity, specificity, precision, etc. determined with a cross-validation or bootstrapping method on test data (see, e.g., [40, 57, 61, 38, 42, 34]). Some other criteria assess the complexity of rules or rule sets such as a rule system size measure (depends on the number of rules), a computational complexity measure (the CPU run-time needed to evaluate a rule or a rule system), a rule complexity measure (number of attributes considered by a rule), a mean scoring rules measure (the average number of rules that are evaluated to find a conclusion), a fuzzy quality measure (for linguistic terms associated with rules), or the information gain for association rules (see, e.g., [4, 30, 37, 14]). Sometimes, several measures are combined [4, 60]. In [27], interestingness measures for rules are evaluated with regard to the four properties of confirmation, locality, symmetry, and a property termed Ex_1 which assures that conclusively confirmatory rules are assigned a higher interestingness value than non-conclusively confirmatory rules and vice versa. Especially, in [27] weaker forms of the locality property and EX_1 are proposed together with a new interestingness measure that fulfills the weaker forms of those properties. Two further measures are proposed in [26] and compared to the measure from [27] in terms of the properties they fulfill. Also, in [26] two new Bayesian confirmation measures for the evaluation of rule interestingness measures are proposed. An interestingness measure that does not evaluate rules or classifier components but individual samples is presented in [8]. There, a support vector machine is used and samples close to the decision boundary or distant to any previous samples are classified as interesting. In [66], a support vector machine is used to learn interestingness values of Twitter hashtags from data.

Subjective measures consider additional knowledge about the application and information about the data miner such as skills and needs [49, 47]. Examples for subjective measures are novelty [5, 17], usefulness [17], understandability [17], actionability [10, 9], and unexpectedness [13, 35, 59, 48].

Existing measures are based on different techniques to represent information about the human users and they depend on the form of knowledge representation. Often, Bayesian networks, fuzzy classifiers, or association rules are addressed in related work.

Related work can also be found in the field of *recommender systems* (e.g., in a content-based approach or a collaborative filtering approach) [2, 31].

In this article, we focus on *objective* measures beyond existing measures for validity (e.g., precision, recall, etc.). We also focus on measures for classification rules and rule sets that are not "crisp", but consider samples and knowledge that are "uncertain". More specifically, we consider probabilistic classifiers and stay within a probabilistic methodological framework to assess the rules contained in these classifiers. In Section 5 we will also address the question of whether these measures could be transferred to other kinds of classifiers, too.

3 Methodological Foundations

In this section we will first present the generative classifier paradigm. A *generative* classifier aims at modeling the processes underlying the "generation" of the data [6]. It is termed to be "generative" because if these processes are modeled perfectly, a generative model could be used to generate artificial data with exactly the same characteristics as the real data. In contrast, discriminative classifiers aim at finding the optimal decision boundary directly. Today, these two approaches are often combined to exploit their respective advantages. Here, we use probabilistic techniques for our generative classifiers. In a second part of the section, we will introduce our new measures for knowledge understanding.

3.1 Probabilistic Classifier CMM

The classifiers we are using here are probabilistic classifiers. That is, for a given, specific D-dimensional input sample $\mathbf{x}'$ we want to compute the posterior distribution $p(c|\mathbf{x}')$, i.e., the probabilities for class membership (with classes $c \in \mathcal{C}$) given the input sample $\mathbf{x}'$. To minimize the risk of classification errors we may then select the class with the highest posterior probability (cf. the principle of *winner-takes-all*). Generally, the posterior distribution $p(c|\mathbf{x})$ can be determined by (cf. [22])

$$p(c|\mathbf{x}) = \frac{p(\mathbf{x}|c)p(c)}{p(\mathbf{x})} = \frac{p(c)\sum_{i\in\mathcal{I}} p(\mathbf{x}|c,i)p(i|c)}{p(\mathbf{x})}, \tag{1}$$

where $p(c)$ is a multinomial distribution with parameters that are termed *class priors* and conditional densities $p(\mathbf{x}|c,i)$ that are called *components*. $\mathcal{I}$ is the overall set of components in this model. In an approach with separate sets of components for the different classes (i.e., we uniquely assign the components to classes, $p(c|i) \in \{0,1\}$), we can conclude that we only have to sum up over all components assigned to a certain class to determine the class posteriors:

$$p(c|\mathbf{x}) = \frac{\sum_{i\in\mathcal{I}_c} p(\mathbf{x}|i_c)p(i_c)}{p(\mathbf{x})}, \tag{2}$$

where $\mathcal{I}_c$ (with $\mathcal{I}_c \subset \mathcal{I}$) is the set of components assigned to class $c \in \mathcal{C}$. Note that

$$p(\mathbf{x}) = \sum_{c\in\mathcal{C}} \sum_{i_c\in\mathcal{I}_c} p(\mathbf{x}|i_c)p(i_c) = \sum_{i\in\mathcal{I}} p(\mathbf{x}|i)p(i). \tag{3}$$

The parameters π_i of the multinomial distribution $p(i)$ are called *mixing coefficients*.

Altogether, we have a classifier consisting of a linear combination of components, where each component is described by a distribution $p(\mathbf{x}|c,i)$. To keep the notation uncluttered, a specific component is identified by a single index $i \in \mathcal{I}$ in the following (i.e., $p(\mathbf{x}|i)$) if its class is not relevant. Which kind of density functions can we use for the components? Considering a D-dimensional sample $\mathbf{x}$ it may have D_{cont} continuous (i.e., real-valued) dimensions (attributes, features) and $D_{\text{cat}} = D - D_{\text{cont}}$ categorical ones. Without loss of generality we arrange these dimensions such that

$$\mathbf{x} = (\underbrace{x_1, \ldots, x_{D_{\text{cont}}}}_{\text{continuous}}, \underbrace{\boldsymbol{x}_{D_{\text{cont}}+1}, \ldots, \boldsymbol{x}_D}_{\text{categorical}}). \tag{4}$$

Note that we italicize x when we refer to single dimensions. The continuous part of this vector $\mathbf{x}^{\text{cont}} = (x_1, \ldots, x_{D_{\text{cont}}})$ with $x_d \in \mathbb{R}$ for all $d \in \{1, \ldots, D_{\text{cont}}\}$ is modeled with a multivariate *normal* (i.e., Gaussian) distribution with center $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. That is, with $\det(\cdot)$ denoting the determinant of a matrix we use the model

$$\mathcal{N}(\mathbf{x}^{\text{cont}}|\boldsymbol{\mu},\boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{\frac{D_{\text{cont}}}{2}}\det(\boldsymbol{\Sigma})^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\mathbf{x}^{\text{cont}}-\boldsymbol{\mu})^{\mathrm{T}}\boldsymbol{\Sigma}^{-1}(\mathbf{x}^{\text{cont}}-\boldsymbol{\mu})\right) \tag{5}$$

For many practical applications, the use of Gaussian components or Gaussian mixture models (GMM) can be motivated by the generalized *central limit theorem* (cf., e.g., [15]) which roughly states that the sum of independent samples from any distribution with finite mean and variance converges to a normal distribution as the sample size goes to infinity. Moreover, any continuous distribution can be approximated arbitrarily well by a finite mixture of normal densities [45].

For categorical dimensions we use a 1-of-K_d coding scheme where K_d is the number of possible categories of attribute $\boldsymbol{x}_d$ ($d \in \{D_{\text{cont}}+1, \ldots, D\}$). The value of such an attribute is represented by a vector $\boldsymbol{x}_d = (x_{d_1}, \ldots, x_{d_{K_d}})$ with $x_{d_k} = 1$ if $\boldsymbol{x}_d$ belongs to category k and $x_{d_k} = 0$ otherwise.

The classifier models categorical dimensions by means of a special case of *multinomial* distributions. That is, for an input dimension (attribute) $\boldsymbol{x}_d \in \{\boldsymbol{x}_{D_{\text{cont}}+1}, \ldots, \boldsymbol{x}_D\}$ we use

$$\mathcal{M}(\boldsymbol{x}_d|\boldsymbol{\delta}_d) = \prod_{k=1}^{K_d} (\delta_{d_k})^{x_{d_k}} \tag{6}$$

with parameters $\boldsymbol{\delta}_d = (\delta_{d_1}, \ldots, \delta_{d_{K_d}})$ and restrictions $\delta_{d_k} \geq 0$ and $\sum_{k=1}^{K_d} \delta_{d_k} = 1$.

We assume that the categorical dimensions are mutually independent and that there are no dependencies between the categorical and the continuous dimensions. Then, the component densities $p(\mathbf{x}|i)$ are defined by

$$p(\mathbf{x}|i) = \mathcal{N}(\mathbf{x}^{\text{cont}}|\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i) \cdot \prod_{d=D_{\text{cont}}+1}^{D} \mathcal{M}(\boldsymbol{x}_d|\boldsymbol{\delta}_{d_i}). \tag{7}$$

With this approach it is possible to model multivariate categorical data sets arbitrarily well despite any independence assumption concerning the categorical variables.

Other dimensions (i.e., other feature types that are not continuous or categorical) can be handled by relying on other (hybrid) distributions, which are not considered yet. Beside that, for some feature types conversions under certain assumptions are possible. For instance, rational ($\mathbb{Q}$) and integer ($\mathbb{N}$) dimensions might be interpreted as nearly continuous under appropriate conditions (e.g., if their support set is large enough). Other attributes (e.g., integers with few values actually occurring) can be categorized as preprocessing step before the CMM is trained.

In Section 3.1.2 we compare CMM to some other classifier paradigms regarding classification accuracy. In general, these probabilistic generative classifiers offer some other interesting features: risk minimizing cost functions can easily be combined with probabilistic outputs, class priors can be compensated, different models can easily be combined, or anomaly detection techniques can be defined [6, 18].

3.1.1 Training of CMM

How can the various parameters of the classifier be determined? For a given training set $\mathbf{X}$ with N samples $\mathbf{x}_n$ it is assumed that the $\mathbf{x}_n$ are independent and identically distributed (i.i.d.). First, $\mathbf{X}$ is split into C subsets $\mathbf{X}_c$, each containing all samples of the corresponding class c, i.e.,

$$\mathbf{X}_c = \{\mathbf{x}_n|\mathbf{x}_n \text{ belongs to class } c\}. \tag{8}$$

For each $\mathbf{X}_c$, a mixture model is trained. Here, we perform the parameter estimation by means of a technique called *variational Bayesian inference* (VI) which realizes the Bayesian idea of regarding the model parameters as random variables whose distributions must be trained [22]. This approach has two important advantages over other methods such as a standard expectation maximization (EM) approach. First, the estimation process is more robust, i.e., it avoids "collapsing" components, so-called singularities when the variance in one or more directions vanishes. Second, VI optimizes the number of components on its own by pruning irrelevant components, i.e., those that are not considered to notably contribute to the overall density (which is reflected by a reasonably small mixing coefficient). Therefore we start the training with a relatively large number of components and rely on VI to automatically reduce the number of components to a sufficient number. That is, the number of model components is not a critical parameter that has to be considered by a user. For a more detailed discussion on Bayesian inference, and, particularly, VI see [6]. More details concerning the training algorithm can be found in [22].

At this point, we have obtained parameter estimates for the $p(\mathbf{x}|c, i)$ and $p(i|c)$, cf. Eq. (1). The parameters for the class priors $p(c)$ are estimated with

$$\gamma_c = \frac{|\mathbf{X}_c|}{|\mathbf{X}|} \tag{9}$$

where $|\cdot|$ denotes the cardinality of a set.

It should be noted that most of the measures defined in this article are independent from the specific training technique, e.g., VI or EM.

3.1.2 Classification Performance of CMM Compared to 1NN, SVM, and Decision Trees

To confirm that CMM are comparable to other classification paradigms regarding their classification accuracy and, thus, sufficiently meaningful to define measures for knowledge understanding, we evaluate their performance on 21 benchmark data sets: phoneme, satimage (real-world data from the UCL Machine Learning Group [64]); australian, credit_a, credit_g, ecoli, glass, heart, iris, pendigits, pima, quality, seeds, segment, vehicle, vowel, wine, yeast (real-world data from the UCI Machine Learning Repository [3]); clouds (artificial data from the UCI Machine Learning Repository [3]); ripley (artificial data proposed in [56]); and two_moons (own, artificial data set). Table 1 contains general information about all data sets, e.g., number of samples, feature (attribute) types, and class distribution.

Table 1: General information about 21 benchmark data sets.

Data Set	Number of Samples	Continuous Attributes	Categorical Attributes	Number of Classes	Class Distribution (in %)
australian	690	6	8	2	55.5,44.5
clouds	5000	2	–	2	52.2,50.0
credit_a	690	6	9	2	44.5,55.5
credit_g	1000	7	13	2	70.0,30.0
ecoli	336	7	–	8	42.6,22.9,15.5,10.4,5.9,1.5,0.6,0.6
glass	214	9	–	6	32.7,35.5,7.9,6.1,4.2,13.6
heart	270	6	7	2	44.4,55.6
iris	150	4	–	3	33.3,33.3,33.3
pendigits	10992	16	–	10	10.4,10.4,10.4,9.6,10.4,9.6,9.6,9.6,10.4,9.6,9.6
phoneme	5404	5	–	2	70.7,29.3
pima	768	–	8	2	65.0,35.0
quality	4898	13	–	7	0.4,3.3,29.7,44.9,17.9,3.7,0.1
ripley	1250	2	–	2	50.0,50.0
satimage	6345	5	–	6	24.1,11.1,20.3,9.7,11.1,23.7
seeds	210	7	–	3	33.3,33.3,33.3
segment	2310	19	–	7	14.3,14.3,14.3,14.3,14.3,14.3,14.3
two_moons	14977	2	–	2	49.1,50.9
vehicle	846	18	–	4	23.5,25.7,25.8,25.0
vowel	528	10	–	11	9.1,9.1,9.1,9.1,9.1,9.1,9.1,9.1,9.1,9.1,9.1
wine	178	13	–	3	33.1,39.8,26.9
yeast	1484	8	–	10	16.4,28.1,31.2,2.9,2.3,3.4,10.1,2.0,1.3,0.3

With the training algorithm Sequential Minimal Optimization (SMO) [50], Support Vector Machines (SVM) are one of the most widely used discriminative approaches for pattern recognition tasks. Consequently, we train SVM with the frequently used libsvm library using RBF (or Gaussian) kernels [11]. Another common paradigm are decision trees (DT). Familiar training algorithms are ID3 or its successor C4.5 [51], for instance. Here we use the latter to train DT. Among the advantages of a DT trained with C4.5 is that the tree structure can easily be used to extract human-readable rules (e.g., [39]). The last competitor is the simple one-nearest-neighbor classifier (1NN). Even though it only classifies with respect to the nearest neighbor, it can be shown that for $N \to \infty$ the maximum classification error is at most twice the maximum of a classifier that yields the best possible classification [12].

The performance evaluation for all classifiers is done in the same way. At first, 20% of the samples are taken from each data set to function as an independent test set. On the remaining 80%, a parameter search is done using grid search combined with a 4-fold cross validation. The parameters that perform best over all folds are then selected to determine the performance on the test set. The results on the independent test sets, achieved with the same parameters as for the training sets, are set out in Table 2. We can see that CMM perform comparably well.

Table 2: Classification accuracy (in %) for four classifier paradigms on 21 data sets.

Data Set	CMM	SVM	DT	1NN
australian	**84.1**	79.7	81.9	79.7
clouds	**89.5**	74.9	88.4	84.8
credit_a	**84.8**	83.3	83.3	78.3
credit_g	68.5	**75.0**	69.5	63.5
ecoli	87.9	**90.9**	84.8	81.8
glass	72.1	69.8	67.4	**83.7**
heart	81.5	**85.2**	75.9	70.4
iris	**96.7**	93.3	93.3	93.3
pendigits	**99.5**	99.4	96.2	99.2
phoneme	86.0	78.3	84.1	**89.1**
pima	68.2	**73.4**	**73.4**	66.9
quality	57.1	**58.1**	55.3	50.5
ripley	**92.0**	90.0	90.0	88.0
satimage	**88.1**	87.8	87.4	84.7
seeds	**100.0**	**100.0**	95.2	90.5
two_moons	**100.0**	99.3	99.8	**100.0**
vowel	98.0	82.3	74.7	**98.5**
wine	97.2	97.2	97.2	97.2
yeast	**54.7**	51.7	51.7	50.0

One substantial property of CMM is that the actual number of components in the mixture model has not a strong influence on the decision boundary as long as there is a sufficient number of components in the model. Fig. 2 demonstrates this exemplarily on one of the used data sets (two_moons). Here, a CMM is trained with VI in such a way that the components of the trained models "cover" smaller areas of the input space which leads to models with more components — 6 components in Fig. 2 (a) up to 12 components in Fig. 2 (c). The images show that the decision boundaries of the resulting CMM classifiers are quite similar.

We exploit this property by initiating our VI training with a higher number of initial components and relying on its pruning capabilities.

3.1.3 Rule Extraction From CMM

In some applications it is desirable to extract human-readable rules from a trained classifier. This is possible with our classifier if it is parametrized accordingly. For the moment, we focus on a single component $p(\mathbf{x}|i)$ and omit the identifying index i. Restrictions concerning the covariance matrix $\mathbf{\Sigma}$ or the number of categories are not necessary. However, to extract a rule set from the CMM, the premises for the continuous input dimensions are obtained from their univariate projections on the axes. This implies that the information about dependencies between different dimensions (i.e., given by the covariances) are lost. Thus, we recommend to force the covariance matrices $\mathbf{\Sigma}$ to be diagonal. Then, the multivariate Gaussians $\mathcal{N}(\mathbf{x}^{\text{cont}}|\boldsymbol{\mu}, \mathbf{\Sigma})$ can be split into a product consisting of D_{cont} univariate Gaussians ψ_d with $d = 1, \ldots, D_{\text{cont}}$. In this case, the univariate Gaussians ψ_d are identical to the projections of the corresponding input dimensions. The categorical dimensions can be simplified by considering only categories whose probability is above a certain case-dependent threshold.

A rule set (where each component is represented by exactly one rule) can then be derived from the classifier as follows: the variables are the input variables x_d (components of the D-dimensional input variable $\mathbf{x}$) and the output variable c which represents the class label. The rule premises are realized by conjunctions of the univariate Gaussians ψ_d (i.e., their density $\psi_d(x_d)$) and conjunctions of settings for the categorical dimensions. The settings themselves are written as disjunctions of those categories that have probabilities above a specified threshold. The conclusions (i.e., the class memberships) are given by one class, which is derived from the class affiliation of the component in the premise. The rules have a form which is very similar to that of fuzzy rules, but they have a rather different (i.e., probabilistic) interpretation.

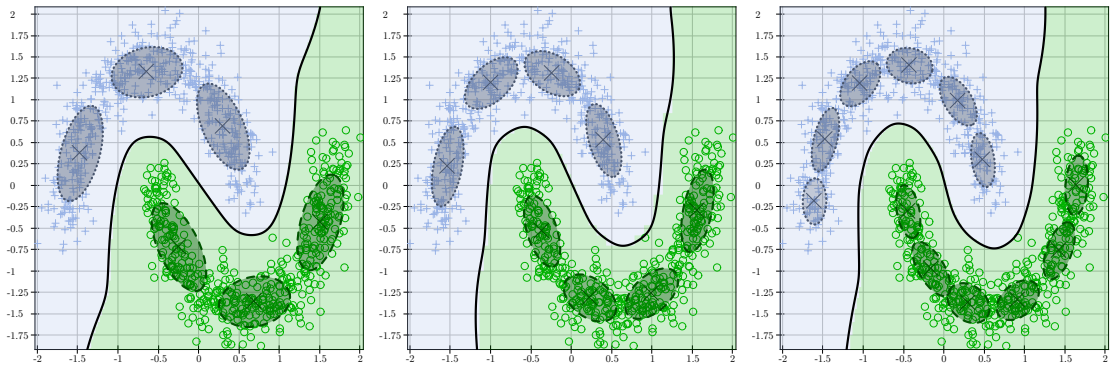

(a) Model with 6 components. (b) Model with 8 components. (c) Model with 12 components.

Figure 2: Different models and classifiers for the artificial *two_ moons* data set, resulting in almost identical decision boundaries. The images show the two-dimensional input spaces of a classifier. Crosses (blue) and circles (green) are samples that originate from processes that can be assigned to two different classes. The probabilistic models being part of the classifiers are depicted by their components (Gaussians) which are represented by ellipses of different shapes (given by the covariance matrices) at positions marked with big ×s (given by the expectations). The thick black lines are decision boundaries which separate the regions assigned to different classes.

Fig. 3 gives an example for a CMM that is used to extract a rule set. The classifier is embedded in a three dimensional input space with two classes: blue crosses (+) and green circles (∘). The first two dimensions x_1 and x_2 are continuous and modeled by three bivariate Gaussians. Their mean values ($\boldsymbol{\mu}_i$) are described by large crosses (+) and the surrounding ellipses are level curves (i.e., surfaces of constant density) of the Gaussians with shapes defined by their covariance matrix $\boldsymbol{\Sigma}_i$. Since the covariance matrices $\boldsymbol{\Sigma}_i$ are only non-zero on their diagonal, all ellipses are axes-oriented. The corresponding projections onto the axes are also illustrated (i.e., the univariate Gaussians). Note, that the projections marked as *high* result from two Gaussians, since the top right component (in green) is covered by the other components (blue) in both continuous dimensions. The third dimension x_3 is categorical with categories A (green), B (blue), and C (red). The distributions of categories are illustrated by the histograms next to every component. Here, only categories with a probability strictly greater than the average $1/K_d$ are considered in order to simplify the resulting rules. Altogether, the following rule set can be extracted from the classifier in Fig. 3:

if x_1 is *low* and x_2 is *high* and (x_3 is A or x_3 is B)
  then c = *blue*
if x_1 is *high* and x_2 is *high* and x_3 is C
  then c = *green*
if x_1 is *high* and x_2 is *low*
  then c = *blue*

Of course, this readability is accomplished at the cost of a limited modeling capability of the classifier (i.e., diagonal covariance matrices and simplified categorical dimensions) and should, thus, only be used if the application requires this kind of human-readable rules.

3.2 Measures for Knowledge Understanding

In the following, we describe seven *objective* measures that can be used to assess the knowledge incorporated in CMM in an objective way. We will use the term rule instead of component only if we wish to explicitly extract human-readable rules from the CMM. In this article, we focus on measures for single components (i.e., rules). Measures for overall classifiers could easily be

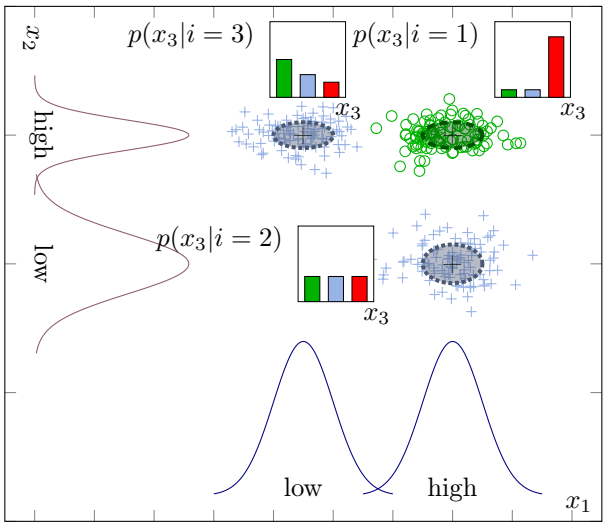

Figure 3: Example of a classifier consisting of three components with two continuous (x_1, x_2) and a single categorical dimension (x_3). Simplified projections of the continuous densities are depicted at the corresponding axes. The distribution of the categorical dimension is shown as histogram attached to the affiliated component.

obtained by averaging measures for components or by considering worst cases etc. Then, it would also be possible to compare different classifiers, for instance. In addition, classical performance measures (e.g., classification error on independent test data) should also be used. If the class a component belongs to is not relevant for a measure, the component is identified by a single index $i \in \{1, \ldots, I\}$ with $I = |\mathcal{I}|$, e.g., $p(\mathbf{x}|i)$. Otherwise, it is explicitly denoted by $p(\mathbf{x}|c, i)$. If sample data are needed to evaluate a measure, we use the training data for that purpose.

3.2.1 Informativeness

A component of the CMM is regarded as being very informative if it is assumed to describe a really distinct kind of process "generating" data. To assess the **informativeness** of a component numerically, we use the *Hellinger* distance $\mathrm{Hel}(p(\mathbf{x}), q(\mathbf{x}))$ of two probability densities $p(\mathbf{x})$ and $q(\mathbf{x})$ (cf. [6]). Compared to other statistical distance measures such as the Kullback-Leibler divergence (cf. Section 3.2.5), the Hellinger distance has the advantage of being bounded between 0 and 1. It is defined by

$$\mathrm{Hel}(p(\mathbf{x}), q(\mathbf{x})) = \sqrt{1 - \mathrm{BC}(p(\mathbf{x}), q(\mathbf{x}))}, \tag{10}$$

where $\mathrm{BC}(p(\mathbf{x}), q(\mathbf{x}))$ denotes the *Bhattacharyya coefficient*:

$$\mathrm{BC}(p(\mathbf{x}), q(\mathbf{x})) = \int \sqrt{p(\mathbf{x}) \cdot q(\mathbf{x})}\, \mathrm{d}x \tag{11}$$

or for discrete distributions with definition range $\mathcal{X}$:

$$\mathrm{BC}(p(\mathbf{x}), q(\mathbf{x})) = \sum_{x \in \mathcal{X}} \sqrt{p(\mathbf{x}) \cdot q(\mathbf{x})} \tag{12}$$

$\mathrm{Hel}(p(\mathbf{x}), q(\mathbf{x}))$ is 0 if $p(\mathbf{x})$ and $q(\mathbf{x})$ are identical and it approaches 1 when $p(\mathbf{x})$ places most of its probability mass in regions where $q(\mathbf{x})$ assigns a probability of nearly zero and vice versa.

Using Fubini's theorem (cf. [25, 1]) and considering the discrete nature of the multinomial distribution, the Bhattacharyya coefficient of two components i and i' as defined in Eq. (7) can be computed by

$$\begin{aligned}\mathrm{BC}(p(\mathbf{x}|i), p(\mathbf{x}|i')) = &\int \sqrt{\mathcal{N}(\mathbf{x}^{\mathrm{cont}}|\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)\mathcal{N}(\mathbf{x}^{\mathrm{cont}}|\boldsymbol{\mu}_{i'}, \boldsymbol{\Sigma}_{i'})}\,\mathrm{d}\mathbf{x}^{\mathrm{cont}} \\ &\cdot \prod_{d=D_{\mathrm{cont}}+1}^{D} \sum_{k=0}^{K_d} \sqrt{\mathcal{M}(\boldsymbol{e}_k|\boldsymbol{\delta}_{d_i})\mathcal{M}(\boldsymbol{e}_k|\boldsymbol{\delta}_{d_{i'}})}\end{aligned} \tag{13}$$

with $\boldsymbol{e}_k$ being the k-th row of the $K_d \times K_d$ identity matrix (i.e., we are iterating over all K_d possible categories of dimension d). The integral can be solved analytically for Gaussians yielding

$$\begin{aligned}&\int \sqrt{\mathcal{N}(\mathbf{x}^{\mathrm{cont}}|\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)\mathcal{N}(\mathbf{x}^{\mathrm{cont}}|\boldsymbol{\mu}_{i'}, \boldsymbol{\Sigma}_{i'})}\,\mathrm{d}\mathbf{x}^{\mathrm{cont}} \\ &= \exp\left(-\frac{1}{8}(\boldsymbol{\mu}_i - \boldsymbol{\mu}_{i'})^{\mathrm{T}}\left(\frac{\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_{i'}}{2}\right)^{-1}(\boldsymbol{\mu}_i - \boldsymbol{\mu}_{i'})\right) \cdot \frac{\sqrt[4]{\det(\boldsymbol{\Sigma}_i)\cdot\det(\boldsymbol{\Sigma}_{i'})}}{\sqrt{\det(\frac{\boldsymbol{\Sigma}_i+\boldsymbol{\Sigma}_{i'}}{2})}}.\end{aligned} \tag{14}$$

The informativeness of a component i is then determined by its Hellinger distance calculated with respect to the "closest" component i' ($i' \neq i$) contained in the CMM:

$$\mathbf{info}(i) := \min_{i' \neq i}\left(\mathrm{Hel}\left(p(\mathbf{x}|i'), p(\mathbf{x}|i)\right)\right). \tag{15}$$

To assess the informativeness of the overall classifier a weighted average of the informativeness values of all components may be used. The weights can be determined depending on the respective mixing coefficients and the class priors.

The run-time complexity required to evaluate the informativeness of one component is

$$\mathcal{O}\left(I \cdot \left(D_{\mathrm{cont}}^3 + D_{\mathrm{cat}} \cdot \max_{d \in \{1,\ldots,D_{\mathrm{cont}}\}}\{K_d\}\right)\right) \tag{16}$$

since for each other component in the classifier we have to compute the determinant of its covariance matrix and iterate over all categories of its categorical dimensions.

3.2.2 Uniqueness

The knowledge contained in the components of a CMM should be unambiguous. This is measured by the **uniqueness** of a component i which reflects to which degree samples belonging to different classes are covered by that component. Let $\rho_i(\mathbf{x}_n)$ denote the *responsibility* of component i for the generation of sample $\mathbf{x}_n$, i.e.,

$$\rho_i(\mathbf{x}_n) := \frac{p(\mathbf{x}_n|i)p(i)}{p(\mathbf{x}_n)}. \tag{17}$$

Then, we define the uniqueness of component i by

$$\mathbf{uniq}(i) := \frac{\sum_{\mathbf{x}_n \in \mathbf{X}_c} \rho_i(\mathbf{x}_n)}{\sum_{\mathbf{x}_n \in \mathbf{X}} \rho_i(\mathbf{x}_n)}. \tag{18}$$

To evaluate the uniqueness of a whole classifier we may, e.g., compute the weighted average of the individual components' uniqueness values (e.g., using the mixing coefficients as weights).

The run-time complexity required to evaluate the uniqueness of a component is $\mathcal{O}(N \cdot I \cdot (D_{\mathrm{cont}}^3 + D_{\mathrm{cat}}))$ since we have to evaluate the density of each sample for the whole CMM to get the responsibilities which involves a matrix multiplication and the iteration over all categorical dimensions.

3.2.3 Importance

The **importance** of a component measures the relative weight of a component within the classifier. In general, either a small or a large number of components may be regarded as "important", depending on a concrete application. Here, a component i is regarded as very important if its mixing coefficient π_i is far above the average mixing coefficient $\overline{\pi} = \frac{1}{I}$. To scale the importance of a component to the interval $[0, 1]$ we additionally use a boundary function that is comprised of two linear functions. One projects all mixing coefficients that are smaller than the average to the interval $[0, 0.5]$ and the other one maps all mixing coefficients that are larger than the average to $[0.5, 1]$. The importance of component i is then computed by

$$\textbf{impo}(i) := \begin{cases} \frac{\pi_i}{2\overline{\pi}}, & \pi_i \leq \overline{\pi} \\ \frac{1-\pi_i}{2(\overline{\pi}-1)} + 1, & \pi_i > \overline{\pi} \end{cases}. \tag{19}$$

Again, to evaluate the importance with regard to a whole classifier we may, e.g., use a weighted average of the importance values of the contained components.

The run-time complexity required to evaluate the importance of one component is $\mathcal{O}(1)$ since it just involves constant time computations with the mixing coefficient of the component.

3.2.4 Discrimination

The **discrimination** measure evaluates the influence of a component i on the decision boundary—and, thus, on the classification performance—of the overall classifier. To calculate the discrimination of component i we create a second CMM by removing i from the original CMM and re-normalizing the mixing coefficients of the remaining components. Then, we compare the achieved classification error on training data (or on test data where available) of the original CMM ($\mathcal{E}_{\text{with}}$) to the classification error of the CMM without component i ($\mathcal{E}_{\text{without}}$):

$$\textbf{disc}(i) := \mathcal{E}_{\text{without}} - \mathcal{E}_{\text{with}}. \tag{20}$$

If required by the concrete application (e.g., in some medical applications false positives are acceptable whereas false negatives could be fatal), it is also possible to use more detailed measures such as sensitivity, specificity, or precision to assess the discrimination capability of a component. We also may consider the class priors.

The run-time complexity required to evaluate the discrimination of one component is $\mathcal{O}(N \cdot I \cdot |\mathcal{C}| \cdot D_{\text{cont}}^3)$ since for each training sample we have to evaluate the density of each component for each class.

3.2.5 Representativity

The performance of a generative classifier also depends on how well it models the data. This kind of fitness is determined by the continuous dimensions only where we explicitly assume that the data distribution can be modeled by a mixture of Gaussians. Since for categorical data it is always possible to find a distribution that perfectly models the data (cf. determination of a histogram), the **representativity** measure only considers the continuous dimensions $\mathbf{x}^{\text{cont}}$. As the true underlying distribution $q(\mathbf{x}^{\text{cont}})$ is unknown for real-world data sets, it must be approximated with a non-parametric density estimation technique, e.g., a standard *Parzen window* density estimator:

$$q(\mathbf{x}^{\text{cont}}) = \frac{1}{N} \sum_{\mathbf{x}_n \in \mathbf{X}} \frac{1}{(2\pi h^2)^{\frac{D_{\text{cont}}}{2}}} \cdot \exp\left(-\frac{\|\mathbf{x}^{\text{cont}} - \mathbf{x}_n^{\text{cont}}\|^2}{2 \cdot h^2} \right). \tag{21}$$

Here, h is a user-defined parameter. Suitable values of h depend on the data set $\mathbf{X}$ [6], but there are a number of heuristics to estimate h. In [7], for instance, h is set to the average distance of

the ten nearest neighbors for each sample, averaged over the whole data set. This non-parametric approach makes no assumptions about the functional form of the underlying distribution.

The calculation of the representativity of the classifier is based on a divergence measure. Here, we rely on a variant of the Kullback-Leibler divergence $\mathrm{KL}(p_1(\mathbf{x})||p_2(\mathbf{x}))$ which for two distributions $p_1(\mathbf{x})$ and $p_2(\mathbf{x})$ is defined as

$$\mathrm{KL}(p_1(\mathbf{x})||p_2(\mathbf{x})) = -\int p_1(\mathbf{x}) \ln \frac{p_2(\mathbf{x})}{p_1(\mathbf{x})} \,\mathrm{d}\mathbf{x}. \tag{22}$$

Other divergence measures, for example, Jensen–Shannon divergence, could also be used. Since the measure is not symmetric, i.e., $\mathrm{KL}(p_1(\mathbf{x})||p_2(\mathbf{x})) \neq \mathrm{KL}(p_2(\mathbf{x})||p_1(\mathbf{x}))$, we use a variant which we denote $\mathrm{KL}_2(p_1(\mathbf{x}), p_2(\mathbf{x}))$. It is given by:

$$\mathrm{KL}_2(p_1(\mathbf{x}), p_2(\mathbf{x})) = \frac{1}{2}\big(\mathrm{KL}(p_1(\mathbf{x})||p_2(\mathbf{x})) + \mathrm{KL}(p_2(\mathbf{x})||p_1(\mathbf{x}))\big). \tag{23}$$

This measure always takes values greater than or equal to 0 and only vanishes if $p_1(\mathbf{x})$ and $p_2(\mathbf{x})$ are identical. It would also be possible to use the Hellinger distance $\mathrm{Hel}(p(\mathbf{x}^{\mathrm{cont}}), q(\mathbf{x}^{\mathrm{cont}}))$, cf. Eq. (10), to measure the distance between the true data distribution $q(\mathbf{x}^{\mathrm{cont}})$ and the model $p(\mathbf{x}^{\mathrm{cont}})$, cf. Eq. (3) and (5). However, since this measure is restricted to the unit interval, errors of the approximation in Eq. (24) may result in values close to 1 and, thus, the difference between the models with and without a certain component is typically close to zero. This effect is alleviated by using KL_2 which has no upper bound and, thus, the influence of a component on the representativity of the model can be quantified with reasonable precision even in the presence of approximation errors.

For the given distribution types Eq. (23) cannot be solved analytically. However, given a dataset $\mathbf{X}$ whose elements are distributed according to $p_1(\mathbf{x})$, the KL_2 divergence can be approximated as follows (cf. the concept of *importance sampling*)

$$\widehat{\mathrm{KL}_2}(p_1(\mathbf{x}), p_2(\mathbf{x})) \approx \frac{1}{2N}\left(\sum_{\mathbf{x}_n \in \mathbf{X}} \ln \frac{p_1(\mathbf{x}_n)}{p_2(\mathbf{x}_n)} + \sum_{\mathbf{x}_n \in \mathbf{X}} \frac{p_2(\mathbf{x}_n)}{p_1(\mathbf{x}_n)} \ln \frac{p_2(\mathbf{x}_n)}{p_1(\mathbf{x}_n)}\right). \tag{24}$$

Representativity evaluates the influence of a component on the "goodness of fit" of the model regarding the data distribution. To calculate the representativity of component i we again create a second CMM without i as described for the discrimination measure. Then, we compare the symmetric $\widehat{\mathrm{KL}_2}$ distance of the CMM with ($p_{\mathrm{with}}(\mathbf{x}^{\mathrm{cont}})$) and without ($p_{\mathrm{without}}(\mathbf{x}^{\mathrm{cont}})$) component i:

$$\mathbf{repr}(i) := \widehat{\mathrm{KL}_2}(p_{\mathrm{without}}(\mathbf{x}^{\mathrm{cont}}), q(\mathbf{x}^{\mathrm{cont}})) - \widehat{\mathrm{KL}_2}(p_{\mathrm{with}}(\mathbf{x}^{\mathrm{cont}}), q(\mathbf{x}^{\mathrm{cont}})). \tag{25}$$

To assess the representativity of a whole classifier we could again use a weighted average of the representatvity values of the contained components. Alternatively we could directly use

$$\widehat{\mathrm{KL}_2}(p_{\mathrm{with}}(\mathbf{x}^{\mathrm{cont}}), q(\mathbf{x}^{\mathrm{cont}})) \tag{26}$$

as an assessment of the representativity of the whole classifier.

The run-time complexity required to evaluate the representativity of one component is $\mathcal{O}(N \cdot I \cdot D_{\mathrm{cont}}^3)$ since we have to evaluate the density of each sample which involves a matrix multiplication.

3.2.6 Uncertainty

The parameters of the CMM are estimated in a Bayesian fashion (cf. Section 3.1), i.e., they are regarded as random variables whose distributions must be determined from sample data. For the parameters of the categorical dimensions $\boldsymbol{\delta}_{d_i}$ the corresponding distributions are *Dirichlet*

distributions. The centers $\boldsymbol{\mu}_i$ and covariance matrices $\boldsymbol{\Sigma}_i$ of the continuous dimensions are modeled with *Gaussian-Inverse-Wishart* distributions [6].

To quantify the **uncertainty** of these parameter estimates we use the *entropy* $\mathcal{H}$ [6]. The entropy of a continuous random variable $\mathbf{x}$ with density $p(\mathbf{x})$ is

$$\mathcal{H}[\mathbf{x}] = -\int p(\mathbf{x}) \ln p(\mathbf{x}) \, \mathrm{d}\mathbf{x}. \tag{27}$$

The more "concentrated" the probability mass of the distribution $p(\mathbf{x})$ is, i.e., the more certain the parameter estimate is, the lower is the entropy $\mathcal{H}[\mathbf{x}]$. Its value is unbounded and can even be negative for continuous variables.

It is also possible to measure the uncertainty for every model parameter individually. In this work, however, we want to calculate an aggregated value quantifying the uncertainty of the estimation of component i. Thus, we sum up the entropies of the corresponding parameter distributions:

$$\mathsf{unce}(i) := \mathcal{H}[\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i] + \sum_{d=D_{\text{cont}}+1}^{D} \mathcal{H}[\boldsymbol{\delta}_{d_i}] \tag{28}$$

This summation of entropies naturally arises from the joint parameter distribution due to the assumption that the continuous dimensions are independent from the categorical ones and the categorical dimensions are mutually independent. Since the absolute value of the entropy values of the categorical dimensions depends on the number of categories and the sum of entropies of all categorical dimensions depends on the number of dimensions in applications where those numbers are very different it may be desirable to weight the summands in Eq. (28) with factors depending on the number of categories per dimension and/or the number of categorical dimensions.

Note that we do not consider the mixing coefficients here. There is only one distribution for the whole classifier and, thus, the same entropy value would be added to every component. The entropies of the mentioned distributions are

$$\begin{aligned}\mathcal{H}[\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i] = & -\frac{D_{\text{cont}}}{2} \log \beta_i + \frac{D_{\text{cont}}(D_{\text{cont}}+1)}{4} \log \frac{\pi}{4} - \frac{D_{\text{cont}}}{2} \log(\det(\boldsymbol{W}_i)) + \frac{D_{\text{cont}}(\nu_i+1)}{2} \\ & - \frac{\nu_i + D_{\text{cont}} + 2}{2} \sum_{d=1}^{D_{\text{cont}}} \psi\left(\frac{\nu_i + 1 - d}{2}\right) + \sum_{d=1}^{D_{\text{cont}}} \log \Gamma\left(\frac{\nu_i + 1 - d}{2}\right)\end{aligned} \tag{29}$$

$$\mathcal{H}[\boldsymbol{\delta}_{d_i}] = -\sum_{k=1}^{K_d} (\epsilon_{d_{k\,i}} - 1)\big(\psi(\epsilon_{d_{k\,i}}) - \psi(\widehat{\epsilon_{d_i}})\big) - \ln C(\boldsymbol{\epsilon}_{d_i}), \tag{30}$$

cf. [6, 36], for instance. Here $\boldsymbol{\epsilon}_{d_i} = (\epsilon_{d_{1\,i}}, \ldots, \epsilon_{d_{K_d\,i}})$ is the parameter vector of the second-order Dirichlet distribution for categorical dimension d_i and

$$\widehat{\epsilon_{d_i}} = \sum_{k=1}^{K_d} \epsilon_{d_{k\,i}}, \tag{31}$$

$$C(\boldsymbol{\epsilon}_{d_i}) = \frac{\Gamma(\widehat{\epsilon_{d_i}})}{\Gamma(\epsilon_{d_{1\,i}}) \ldots \Gamma(\epsilon_{d_{K_d\,i}})}. \tag{32}$$

$\Gamma(\cdot)$ is the gamma function defined by

$$\Gamma(x) = \int_0^\infty t^{x-1} \exp(-t) \, \mathrm{d}t, \tag{33}$$

and $\psi(\cdot)$ is the digamma function defined by (cf. [6])

$$\psi(x) = \frac{\mathrm{d}}{\mathrm{d}x} \ln \Gamma(x). \tag{34}$$

In contrast to our previous measures this measure uses the second-order distributions which arise during VI training (cf. Section 3.1.1) as distributions over the actual parameters of our CMM. More details on training models using second-order distributions can be found in [6, 22, 20].

To evaluate the uncertainty of a whole classifier we could again compute a weighted average of the uncertainty values of the contained components using the mixing coefficients. Additionally, we might take the uncertainty of the second-order distribution which models the mixing coefficients into account. This uncertainty is modeled by the variances of these distributions.

The run-time complexity required to evaluate the uncertainty of one component is

$$\mathcal{O}\left(D_{\text{cont}}^3 + D_{\text{cat}} \cdot \max_{d \in \{1,\ldots,D_{\text{cat}}\}} \{K_d\}\right), \tag{35}$$

since we have to compute the determinant of a matrix for the continuous dimensions and sum over the categories of all categorical dimensions.

3.2.7 Distinguishability

For the rule set represented by the classifier to be easily *understandable* for a human expert each two rules should be easily distinguishable. Since categorical dimensions are easily distinguishable by humans we only consider the continuous dimensions here. To measure the distinguishability of two components i and i' for dimension $d \in \{1, \ldots, D_{\text{cont}}\}$ we first project those components onto d. This yields two univariate Gaussians φ_{d_i} and $\varphi_{d_{i'}}$. In order to restrict our **distinguishability** measure to the unit interval we omit the normalizing coefficients of the projected Gaussians. This also makes sure that strongly overlapping Gaussians which are not easily distinguishable for a human are assigned a low distinguishability value. To assess the distinguishability of the two components i and i' with regard to dimension d we now use the intersection point of the two projected Gaussians that lies between their centers which results in

$$\textbf{dsng}_d(i, i') := 1 - \exp\left(-\frac{(\mu_{d_i} - \mu_{d_{i'}})^2}{2(\sigma_{d_i} + \sigma_{d_{i'}})^2}\right) \tag{36}$$

with $\textbf{dsng}(i, i') \in (0, 1]$. We now aggregate those values over all components i' different from i and over all continuous dimensions d to get the overall distinguishability for component i as

$$\textbf{dsng}(i) = \min_{d \in \{1,\ldots,D_{\text{cont}}\}} \left\{ \min_{\substack{i \in \mathcal{I} \\ i \neq i'}} \{\textbf{dsng}_d(i, i')\} \right\}. \tag{37}$$

Values higher than 0.1, for example, could be regarded as desirable, depending on the application. Fig. 4 shows an example of the assessment of the distinguishability of two exemplary Gaussians.

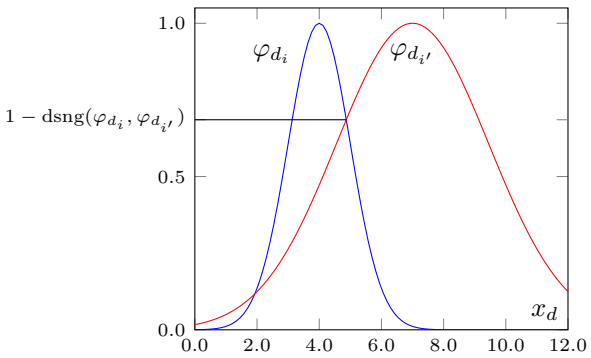

Figure 4: Example of an assessment of the distinguishability of two Gaussians.

The run-time complexity required to evaluate the **distinguishability** of one component is $\mathcal{O}(D_{\text{cont}} \cdot I)$ since for each continuous dimension we have to iterate over all components of the classifier to compute the intersection points of the projected Gaussians.

The distinguishability of a set of rules (i.e. the whole classifier) is defined to be the distinguishability of the rule with the lowest distinguishability. When evaluating the complete rule set two additional factors may be considered regarding the interpretability of the rule set by a human expert.

First, the number I of rules in the rule set should be low because we argue that a classifier with few rules is easier to understand than a classifier with many rules.

Second, the number of different terms τ_d for each dimension d should be low. For a categorical dimension τ_d is given by the number of different categories with non-zero probability forming the disjunctions. To simplify categorical dimensions only categories with a probability above a certain threshold $\omega > 0$ may be considered (cf. Section 3.1.3). Let $\mathcal{K}_{d_i}$ be the set of categories of dimension d of component i. Then we have

$$\tau_d = \left| \bigcup_{i=1}^{I} \mathcal{K}_{d_i} \right|. \tag{38}$$

For a continuous dimension d the number of different univariate Gaussians φ_{d_i} is counted. To decide whether two Gaussians should be regarded as being different or not, we use the Hellinger distance (cf. Eq. (10)) of the two Gaussians. The distance should be clearly below a fixed threshold, such as 0.01, for example, to regard two Gaussians as being identical:

$$\tau_d = \sum_{i=1}^{I} h_{d_i} \tag{39}$$

where

$$h_{d_i} = \begin{cases} 0, & \text{if there is a } \varphi_{d_{i'}} \text{ with } i < i' \text{ and } \text{Hel}(\varphi_{d_i}, \varphi_{d_{i'}}) \leq 0.01 \\ 1, & \text{otherwise} \end{cases}. \tag{40}$$

The threshold can be varied in applications depending on the degree of distinguishability that is desired. To assess a complete rule set we could average the number of different terms for all categorical and continuous dimensions. For the classifier shown in Fig. 3 we get $\tau_1 = 2$, $\tau_2 = 2$, and $\tau_3 = 3$ resulting in $\tau = 2.3$, for example.

4 Case Studies

In this section we investigate the properties of the proposed measures in detail by (1) analyzing correlations between these measures and run-times, and by (2) conducting four case studies that demonstrate how the measures can be used in practical applications. These case studies show how measures can be used in the learning phase of the classifier, to improve the classification performance in an active learning setting, while evaluating the trained classifier before using it on-line, and finally during the application phase of the classifier. These case studies can be seen as illustrative examples; many other ways of using the measures are possible.

4.1 Correlation Analysis and Run-Time

In the first set of experiments we analyze correlations between the seven measures to investigate their dependencies. Additionally, we measure the run-time of all evaluations of our measures required to compute those correlations to get some empirical evidence to back up the theoretical run-times stated in Section 3.2. For this experiment we again use the 21 benchmark data sets.

Table 3: Spearman's correlation coefficients averaged for the seven interestingness measures over 21 benchmark data sets (with standard deviations). Entries which are statistically significant for a significance level of 0.05 are highlighted.

	info	**impo**	**uniq**	**disc**	**repr**	**unce**	**dsng**
info	1.00 ± 0.00	-0.23 ± 0.39	$\mathbf{0.56 \pm 0.38}$	0.07 ± 0.28	-0.14 ± 0.44	0.19 ± 0.44	-0.01 ± 0.37
impo		1.00 ± 0.00	-0.19 ± 0.36	$\mathbf{0.55 \pm 0.27}$	$\mathbf{0.56 \pm 0.28}$	$\mathbf{-0.51 \pm 0.69}$	0.24 ± 0.34
uniq			1.00 ± 0.00	-0.01 ± 0.26	-0.13 ± 0.47	0.24 ± 0.40	0.09 ± 0.34
disc				1.00 ± 0.00	0.34 ± 0.28	-0.18 ± 0.42	0.09 ± 0.33
repr					1.00 ± 0.00	-0.17 ± 0.57	0.18 ± 0.30
unce						1.00 ± 0.00	-0.07 ± 0.44
dsng							1.00 ± 0.00

Table 3 shows Spearmans's rank correlations computed for our seven interestingness measures and averaged over the 21 data sets. Correlations that are statistically significant on a significance level of 0.05 are highlighted with bold typeface in Table 3. Those significant correlations are:

- **Informativeness** is positively correlated to **uniqueness**. This means that isolated components do not cover many samples. Furthermore, components that are located relatively close to other components of a CMM cover a majority of the samples. This may be due to the fact that many of the benchmark data sets consist of real-world data for which the normal distribution assumption may not be fully satisfied. This leads to clusters being modeled by multiple components rather than by one single component. It is also consistent with importance being negatively correlated to uncertainty.

- **Importance** is positively correlated to **discrimination**, i.e., components that cover many samples also have a high impact on the decision boundary of the classifier.

- **Uncertainty** is negatively correlated to **importance** which means that components covering only very few samples yield only very little additional information which expressed through the entropy yields the uncertainty.

- **Representativity** is positively correlated to **importance**. This is due to the fact that components that have a high mixing coefficient in comparison with the other components of the model also have a high influence on the overall density function which is evaluated by the **representativity** measure.

For several combinations of interestingness measures the values of the correlation coefficients are quite. Especially, the **distinguishability** measure has no significant correlation (considering Spearman's correlation coefficient) with any of the other measures. This shows that the **distinguishability** measure evaluates completely different aspects of the components of a CMM classifier than the other measures. For example, a component can have a high **distinguishability** while still having either a small or big influence on the decision boundary as measured by **discrimination**. Apart from the **distinguishability**, the pair of measures with the lowest correlation is **discrimination** and **uniqueness**. If a component has a very high **uniqueness** it mostly "covers" samples of one class and if its **uniqueness** is very low it "covers" samples of several classes. In both cases the **discrimination** of the component, i.e. its influence on the decision boundary, can either be very high or low. This just depends on the position of the component with regard to its neighboring components and the decision boundary and notably not on the **uniqueness** of the component. Altogether, the fact that there are only a few significant correlations between our objective interestingness measures indicates that our measures cover many different aspects of knowledge understanding.

We also measured the empirical run-times it took to compute the results given in Table 3. The results are presented in Table 4 and were obtained using a dedicated Linux machine with an Intel Core i7 2600 CPU running at 3.4 GHz. The run-time measurements are predominantly a rough

Table 4: Comparison of run-times in milliseconds required to compute the results in Table 3. The last columns are the number of samples (N) and dimensions (D) of the data set and components (J) in the final model.

	info	**impo**	**uniq**	**disc**	**repr**	**unce**	**dsng**	N	D	J
australian	198	95	2550	3013	163066	26	437	690	14	19
clouds	29	33	88	170	59940	14	30	5000	2	4
credit_a	228	107	6703	6185	43692	65	99	690	15	26
credit_g	3200	720	203792	216984	388329	407	1104	1000	20	107
ecoli	94	19	280	1020	7039	8	14	336	6	22
glass	207	72	394	970	5892	6	87	214	9	21
heart	5	3	94	70	1398	15	2	270	13	5
iris	7	1	32	86	719	1	26	150	4	7
pendigits	8760	8229	104277	1541128	18217114	3058	6076	10992	16	104
phoneme	844	334	17773	14708	1593644	256	415	5404	5	47
pima	851	190	4526	3993	109913	68	92	768	8	41
quality	10254	3570	294072	652816	18145798	1667	4342	4898	11	182
ripley	6	4	22	34	4214	4	86	1250	2	4
satimage	220	147	1075	4341	535105	41	98	6435	5	26
seeds	1	< 1	4	4	583	6	1	210	7	3
segment	2125	953	34704	136056	3133814	416	4643	2310	18	56
two_moons	482	291	1896	2919	1434223	111	159	14977	2	11
vehicle	5055	1688	32658	66724	615648	337	1735	846	18	60
vowel	11493	988	73744	154656	862679	426	3126	990	10	187
wine	499	192	869	781	8258	10	187	178	13	30
yeast	3638	492	35094	150671	606798	202	489	1484	6	57

guideline to estimate how long it takes to evaluate individual measures in actual applications. In our experimental implementation we optimize many computations and rely on as many caches as possible. For example, inverse covariance matrices required to evaluate the densities are already pre-computed by our training algorithm. Thus, in contrast to what is suggested by the theoretical run-time complexities of our interestingness measures as stated in Section 3.2, the empirical run-times are not dominated by matrix operations but rather by the number of samples in each data set. For small data sets such as seeds, ripley, iris, or heart, the evaluations of all measures require less than a second. In comparison, the pendigits and quality data sets require the highest run-times up to a few hours. With regard to the measures, representativity generally takes the longest to evaluate, followed by discrimination and uniqueness. The evaluation of the other measures is very fast.

4.2 Knowledge Acquisition Phase: Controlling the Training of a Classifier

In this set of experiments we use some of our measures, **importance**, **representativity**, and **uncertainty**, to control the VI based training process of a classifier. The training consists of three steps, an E step, an M step and an additional pruning step, which are consecutively executed until a convergence criterion is met. In the E step, samples are gradually assigned to components of the classifier and in the M step, the parameters of the classifier are updated according to the samples (gradually) assigned to them. In the pruning step, unnecessary components are removed from the classifier. Here, we focus on the pruning step and consider three different pruning methods:

1. **Resp:** This method is similar to the "traditional" pruning method used in VI training: A component is removed from the classifier if the sum of its un-normalized responsibilities (cf. Eq. (17)) is *below* a certain threshold. Often, this threshold is chosen as 1 which means that a component is removed if it is effectively "responsible" for less than one sample of the training data. A threshold of 2.5 means the component is "responsible" for less than 2.5 samples and so on. A threshold of 0 is also possible in which case no pruning is done.

2. **Impo:** Based on the **importance** measure (cf. Section 3.2.3) this pruning method uses the mixing coefficient of a component to decide whether or not to prune it. However, in

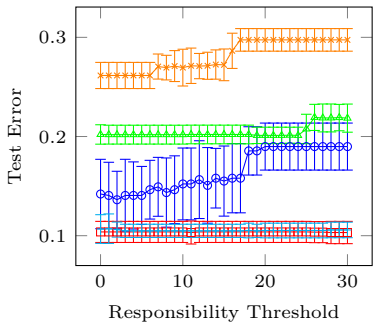

(a) Average classification error on test data with standard deviation for pruning method Resp.

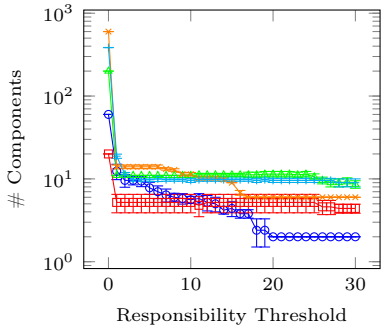

(b) Average number of finally resulting components with standard deviation for pruning method Resp.

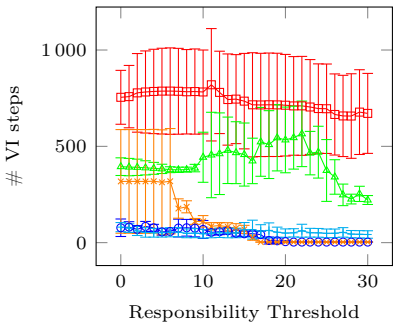

(c) Average number of VI steps with standard deviation for pruning method Resp.

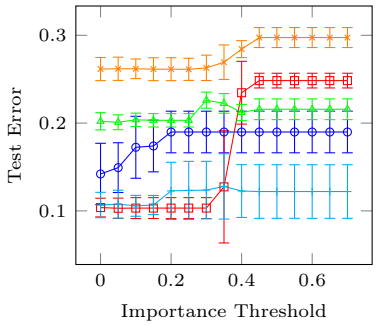

(d) Average classification error on test data with standard deviation for pruning method Impo.

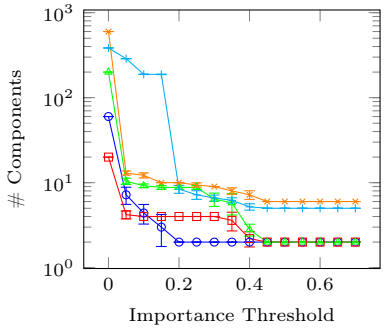

(e) Average number of finally resulting components with standard deviation for pruning method Impo.

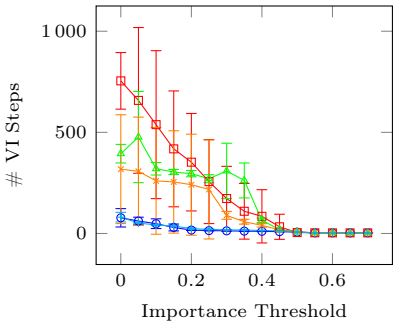

(f) Average number of VI steps with standard deviation for pruning method Impo.

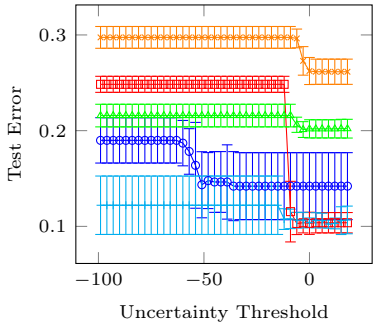

(g) Average classification error on test data with standard deviation for pruning method Unct.

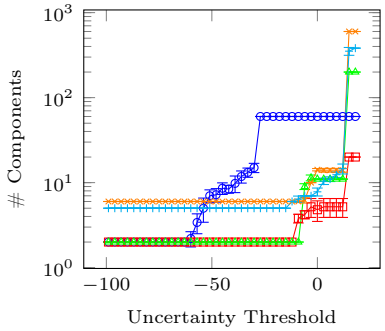

(h) Average number of finally resulting components with standard deviation for pruning method Unct.

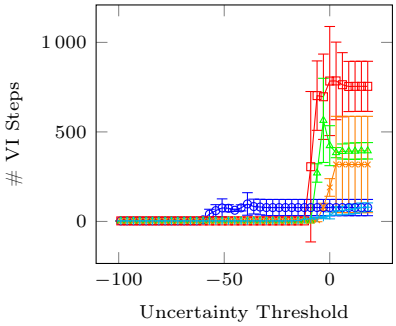

(i) Average number of VI steps with standard deviation for pruning method Unct.

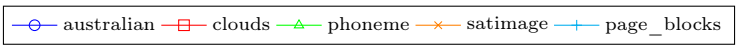

(j) Legend specifying the 5 used benchmark data sets.

Figure 5: Comparison of different pruning methods applied to five benchmark data sets using a 5-fold cross-validation. Data was first reduced to two dimensions using a PCA and then normalized using a z-transform. Note that in images (a) to (f), components are pruned if the measure drops *below* the given threshold (horizontal-axis), whereas in images (g) to (i) components are pruned if the given threshold is *exceeded*.

contrast to Resp not only one component is considered but the mixing coefficient is set in relation to the average mixing coefficient of the classifier (or, in other words, the number of components the classifier currently consists of). Thus, a component is removed if its importance is *below* a certain threshold. Since importance values lie in the interval $[0, 1]$, thresholds in that interval may be considered. Values close to 0 result in slow pruning and values close to 1 lead to fast pruning. A threshold of 0.5 means that all components whose mixing coefficient is below average get pruned.

3. **Unct:** In contrast to Resp and Impo this method does not rely on mixing coefficients or responsibilities but takes a completely different approach. It considers the **uncertainty** of a component (cf. Section 3.2.6) and removes all components whose uncertainty value lies *above* a certain threshold. The uncertainty of a component is unbounded and it is, thus, more difficult to determine a suitable threshold. However, the following examples provide good starting points for choosing a threshold value.

Fig. 5 visualizes classification errors on test data, numbers of resulting classifier components, and numbers of required VI training steps for the three different pruning methods. We used 4 of the benchmark data sets from Section 4.1 for this experiment: australian, clouds, phoneme, and satimage. Additionally, we added the page_blocks data set from the UCI Machine Learning Repository [3] which consists of real-world data. It can be seen that the Resp method yields the best results for low thresholds in the range $[1, 3]$ but still requires a larger number of VI steps than the other methods. In contrast to Resp, the Impo pruning method leads to a faster decrease in the number of required VI steps while resulting in a similar test error and number of components. The threshold used in the **Unct** method is rather sensitive to the type of data set used. The australian data set contains categorical dimensions in addition to continuous ones and this leads to rather different entropy values and, thus, different thresholds. In conclusion, we can state that importance and uncertainty can both be used as alternatives to responsibilities for a pruning step of the VI training algorithm by reducing the number of required VI steps and thus making the training process faster. However, further research into an automatic determination of the pruning threshold and the behavior of the pruning methods on different data sets is necessary.

4.3 Knowledge Analysis Phase: Ranking Components of a Trained Classifier

In this experiment we analyze the components of a trained classifier with our measures in order to help a potential user of the classifier in understanding the components or rules extracted from those components (cf. Section 3.1.3).

We trained a classifier on the phoneme data set from the UCL Machine Learning Group [64] using the restriction to diagonal covariance matrices to enable rule extraction. Fig. 6 shows the resulting classifier. Components are colored according to their class densities and their opacity is dependent on their mixing coefficient. The background of the plot is colored according to the posterior probabilities of the classifier. The black, solid line is the decision boundary.

In Table 5, some of our measures are evaluated for the components depicted in Fig. 6. Other measures were omitted for the sake of brevity.

From the **distinguishability** measure we can deduce that components 3 and 6 are those which can least easily be distinguished. This can be confirmed by looking at Fig. 6 which shows that those components overlap to a great extent. Components 5 and 7 can best be distinguished from others when projected to the coordinate axes which is obvious especially in case of component 5 because it is well separated from all other components.

The second measure evaluated in Table 5 is **discrimination**. Since the discrimination value of component 3 is negative, we know that the classification performance could be increased if we removed this component from the classifier. In this case, the classification error on test data decreases from 21.46% to 21.18%. Also, components 4 and 5 are not very important for the classification decision since they are "dominated" by component 2 as we can see from Fig. 6.

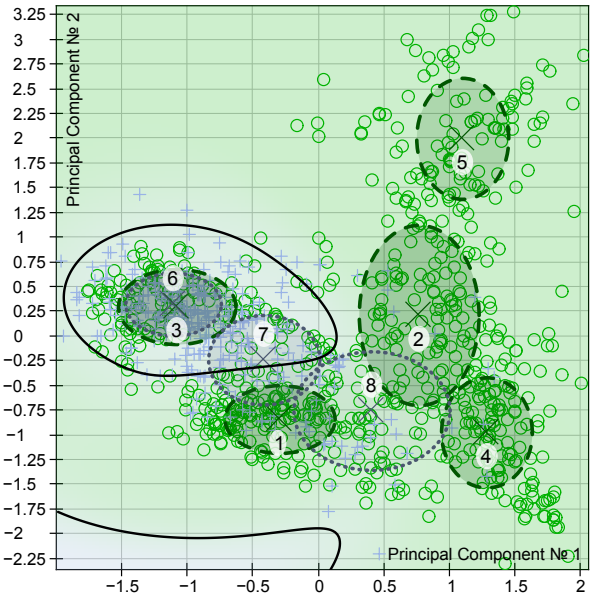

Figure 6: A fraction of the data of the phoneme data set together with a trained classifier. The data were reduced to two dimensions using a PCA. Components of the first class are visualized by dashed green ellipses and numbered below their center. Components of the second class are shown as dotted blue ellipses and are numbered above their center.

The components that are most important for the classification decision are components 1 and 6. Those are the only green colored components in an area otherwise dominated by the blue components 3, 7, and 8.

The informativeness measure info in Table 5 yields components 4 and especially 5 as those with the greatest distance to all other components. In fact, they have nearly no overlap with other components in Fig. 6. The least informative components are components 3 and 6 since they almost totally overlap.

The final measure we evaluated in Table 5 is **uniqueness**. The components with the highest uniqueness are components 2, 4, and in particular 5 which means they mostly cover samples from one class. Components 3 and 8 have the lowest uniqueness which means they cover samples from both classes. Interestingly, component 6 which nearly completely overlaps with component 3 has a higher uniqueness value.

The example shows that the proposed measures can easily be used to analyze a trained classifier. It should be emphasized that the evaluation of a classifier does also work as described in case of higher dimensional input spaces where looking at a two-dimensional visualization would not help very much. Depending on the application, different goals could be achieved by ranking the components of a classifier using our approach. If the aim was to reduce the number of components, for example, one might consider removing component 8 since it received the worst evaluation by four out of seven measures (not all shown here for the sake of brevity): it has the lowest uniqueness, representativity, and importance, and the highest uncertainty.

4.4 Knowledge Analysis in Active Learning

The pool-based active learning (PAL) [41] paradigm repetitively asks users (generally termed as *oracles*) to provide label information for unlabeled data, e.g., in order to train a classifier based on those data. PAL is based on the assumption that unlabeled data can be acquired at no (or low) costs, whereas retrieving label information is very costly. Therefore, at the beginning of a PAL learning cycle a large set of unlabeled data and only a small set of labeled data are available. Based on these data a classifier is trained (here a CMM), which is further on used to

Table 5: Ranking of components based on various measures. Component numbers are given in Fig. 6.

Measure	Value	Rank	Component
dsng	0.0022	1	3
	0.0022	2	6
	0.0037	3	2
	0.0134	4	1
	0.0134	5	8
	0.0498	6	4
	0.0583	7	7
	0.2574	8	5
disc	−0.0028	1	3
	0.0000	2	4
	0.0000	3	5
	0.0028	4	8
	0.0157	5	7
	0.0194	6	2
	0.0722	7	6
	0.1156	8	1

Measure	Value	Rank	Component
info	0.1342	1	3
	0.1342	2	6
	0.5128	3	1
	0.5128	4	7
	0.5221	5	2
	0.5221	6	8
	0.6267	7	4
	0.7401	8	5
uniq	0.2081	1	8
	0.3900	2	3
	0.4515	3	7
	0.6293	4	6
	0.7815	5	1
	0.9074	6	2
	0.9482	7	4
	0.9991	8	5

build a ranking of the unlabeled data based on an estimate how likely is it for a sample to increase the performance of the classifier if it was labeled. A *selection strategy* uses these estimates as decision basis for selecting in each learning cycle the next sample or set of samples that is going to be labeled. Fig. 7 depicts the learning cycle of standard PAL with solid arrows. Typically, a large pool of unlabeled samples (U) and a small set of labeled samples (L) are available at the beginning of PAL, which makes it possible for a classifier $\mathbf{G}$ to be trained. Then, a set S ($S \subset U$) is queried based on a selection strategy $\mathcal{Q}$ and presented to the oracle $\mathcal{O}$ for labeling. The labeled samples ($S_{labeled}$) are added to L and the classifier is updated. As long as a given stopping condition is not met, a new PAL cycle (i.e., learning cycle) is started.

The PAL process starts with a generative (probabilistic) model as described in Section 3.1 which can be trained in an unsupervised way. During the PAL process, more and more labels become available that can be considered to train a CMM. In general, the underlying generative model remains unchanged during the PAL process. In [54], this process is extended with a *transductive learner*, which aims at updating the underlying model using the **uniqueness** measure, which determines how ambiguous the knowledge modeled by each component is. If the uniqueness value determined for a component is smaller than a predefined threshold, the component is considered to be "disputed" between two or more classes. In this case, the samples for which the component is responsible (cf. Eq. 17) are determined, a sample-based classifier called Resp-kNN [52] is trained and used to transductively label the underlying samples. Then, a new CMM is trained and the resulting components *fused* with "non-disputed" components, where necessary(for details of the fusion technique, see [24]). Consequently, the information provided by the uniqueness measure allows for iteratively improving a CMM during a PAL process. For further information regarding PAL with a transductive learner see [54]. This extension of the PAL process is also shown in Fig. 7.

For the case study presented in this section, we embed the probabilistic generative model in a kernel function of a support vector machine (SVM) as described in [55]. In this case study, the **uniqueness** measure is only applied once, after a final PAL cycle, to emphasize on the possible performance improvement.

We have conducted an experiment with the MNIST [46] handwritten digits data set, which consists of a training set of 60 000 and a test set of 10 000 gray-scaled images of handwritten digits (10 classes). We reduced the size of the input dimensions from 784 to 34 (continuous dimensions) by applying a principle component analysis, thus keeping 90% of the total variance. Furthermore, the data has been z-normalized. The PAL process started with an initially labeled set of 40 samples and selected, based on the 4DS [53] selection strategy (corresponds to $\mathcal{Q}$ in Fig. 7) five

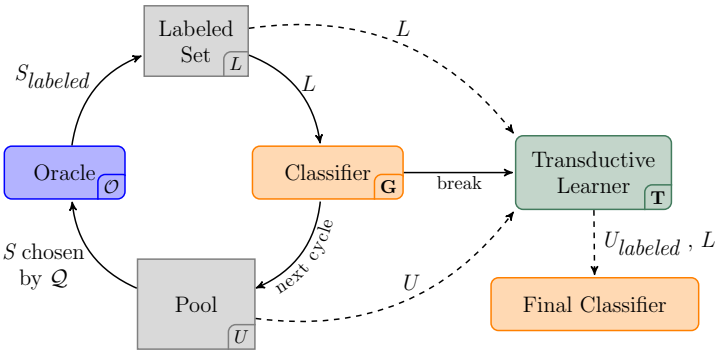

Figure 7: Learning cycle of PAL (solid arrows) with extension (dashed arrows).

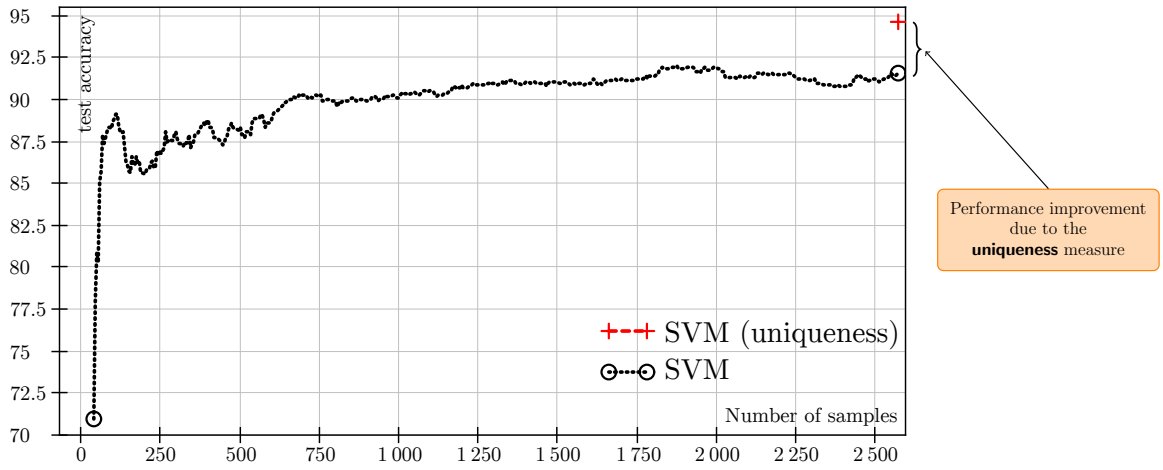

Figure 8: Learning curve (test accuracy versus the number of actively selected samples) of the active training of SVM with generative kernel (based on GMM) for the MNIST data set (magenta). The red + indicates the test accuracy after refining the underlying GMM of the kernel by means of the **uniqueness** measure.

samples in each learning cycle of the PAL process. We decided to actively select 2 575 samples (i.e., the PAL stopped when the size of the labeled set L reached this number). The actively trained classifier (here SVM) had been parametrized only using the training set.

After the PAL process stopped, an adequate number of labeled samples was available. Thus, it was possible to improve the underlying density model by means of the **uniqueness** measure as sketched above. Then, the density model can better exploit the structure information contained in the labeled samples. Fig. 8 illustrates the learning curve of the SVM on the test set. Moreover, it emphasizes that by employing the **uniqueness** measure the classification error is reduced by about a half.

This promising result indicates a possible improvement of the standard PAL process. In regular time intervals (i.e., cycles of the PAL process) the underlying density model should be evaluated by means of the **uniqueness** measure and, if necessary, refined.

4.5 Knowledge Application Phase: Detecting Novel Processes

Our last set of experiments deals with the application phase of a classifier. We consider the task of novelty detection, i.e., the task of detecting the need of new components that model newly occurring processes in the input space which were not known during training time (cf. [43]). Our example is based on intrusion detection data from the 1999 KDD cup [63]. Specifically, we use data of the attack types “neptune”, “smurf”, “ipsweep”, and “satan”. We re-sampled the data to be

able to consider longer time spans. Each of the four data sets we constructed starts with 54 000 samples of background traffic (large variety but without any attacks). Then, an attack phase starts which lasts until time step 198 000. During the attack phase samples that represent the respective attack are present in a ratio of 1:3. Finally, we continue with more samples only drawn from background traffic to observe how the novelty detectors react to the end of an attack. The classifiers were trained with a separate training data set consisting of 5 000 samples of background traffic only. Our goal is to detect the need for a new component that models data originating from an attack to the network. This shows that intrusion detection systems that detect new kinds of attacks at run-time can be build by relying on our proposed interestingness measures.

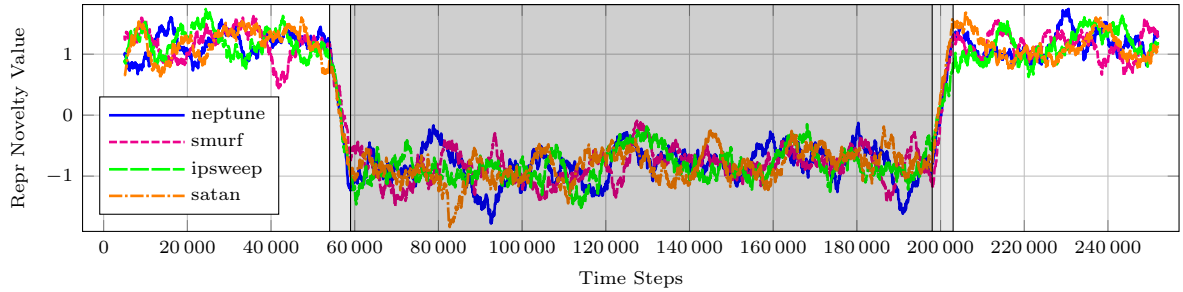

(a) Representativity based novelty detection.

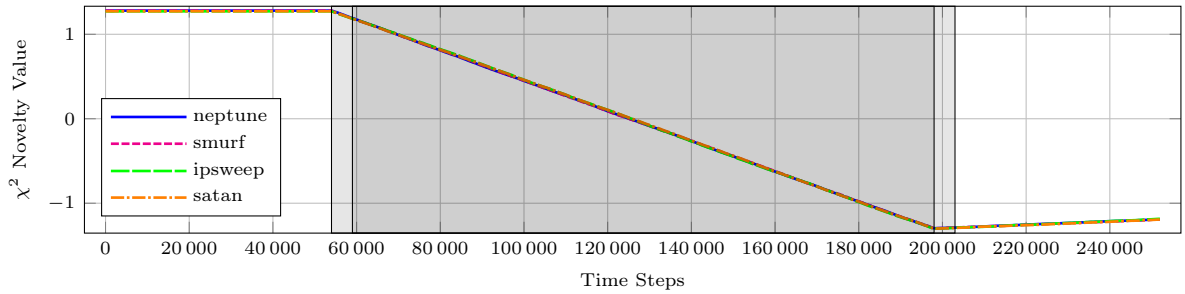

(b) χ^2 novelty detection.

Figure 9: Novelty detection with our representativity measure (top) in comparison to the χ^2 novelty detector as described in [23] (bottom) on 4 intrusion detection data sets. The grayed region marks the time an attack occurred. The light gray area indicates the size of the sliding window used by the representativity measure after the attack started and after it stopped again.

In order to achieve that goal we evaluate the representativity of the overall classifier after each newly observed sample. The Parzen window estimator of the **representativity** measure that is used for the evaluation of repr is computed from a sliding window of the most recent 5 000 samples. In Fig. 9 (a), representativity values of the classifiers are shown. For the sake of visualization, the representativity values were normalized to mean 0 and variance 1. During the initial phase, which contains only background traffic, the representativity values of all classifiers are high. When an attack starts at time step 54 000, the representativity values decrease until the sliding window is filled with data that contains 25% of attack samples. Now, the representativity values stay small until the attack to the network stops at time step 198 000. Then, the representativities of the classifiers rise again to high values, similar to the ones in the initial phase.

For comparison, Fig. 9 (b) contains results obtained by the novelty detection technique called χ^2 novelty which is proposed in [23]. This technique measures whether or not each newly observed sample lies within a certain region around any of the classifier components and accumulates reward or punishment values accordingly. Fig. 9 (b) depicts the accumulated novelty values. For the sake of visualization those values were also normalized to mean 0 and variance 1.

In conclusion we can state that our **representativity** measure can be used to detect the need for new components in a classifier. The raw representativity is more sensitive to novel samples than the χ^2 novelty measure which is more robust, but it is certainly worth to investigate techniques that combine the advantages of both.

5 Conclusion and Outlook

In this article, we presented an approach to support or even automate the process of knowledge understanding, an important data mining step. However, the proposed measures should be seen as a first step into that direction, as not all aspects of real world applications could be considered yet. With knowledge understanding we refer to the task of analyzing the knowledge extracted from data in order to gather meta-knowledge (knowledge about knowledge). This can either be done offline (here referred to as "understanding") or online (here referred to as "experience") by applying the knowledge. Depending on a concrete application, all or only a part of the measures may be useful. Also it may be preferable to assess continuous and categorical dimensions separately.

The measures proposed in this article are objective measures but they have some relation to subjective interestingness measures. For example, **informativeness** is clearly related to *novelty*, a "conventional" *interestingness measure*, since if a rule is quite informative it may also be the case that the user regards this rule as providing a highly novel kind of knowledge. **Uniqueness** can be seen as being related to *understandability*, because a high **uniqueness** value implies that a rule is only associated with a single class, thus making it better understandable. The **importance** and **discrimination** measures are both reflecting the *usefulness* of a component by expressing its relative weight in the classifier and the influence on the decision boundary. Another measure related to *usefulness* is our **uncertainty**: A high value states that the parameters of the process underlying the resulting component cannot be precisely determined, which decreases its *usefulness* of the component. The last measure **representativity**, which expresses how well a component models the data, has a positive effect on the *understandability* of the overall model.

We defined various measures for CMM (classifier based on a mixture model), a generative, probabilistic classifier. In practical applications, this classifier can be used in combination with discriminative classifiers such as support vector machines (SVM). As we have shown in [21], the probabilistic classifiers with continuous input dimensions are *functionally* equivalent to certain kinds of support vector machines (with Gaussian kernels), radial basis function neural networks, fuzzy classifiers (with Gaussian membership functions and sum-prod inference), nonlinear Fisher discrimination techniques, relevance vector machines, or direct kernel machines. Thus, some of the measures defined here could be adapted to these kinds of classifiers, as well. As the semantics underlying the knowledge in these classifiers is very different (cf., e.g., membership degrees in fuzzy systems to probabilities in our generative classifiers), this idea has to be investigated in detail. Another option is to use the probabilistic model contained in our CMM in other classifiers, for instance, as we have already shown in [55], to create a data dependent kernel function for SVM.

In this article, several use cases were defined and investigated to demonstrate the value and applicability of the proposed measures. We have shown that these measures may be used to control the training process of the classifier (for more details on this problem see also [28]), to analyze the knowledge contained in a trained classifier, and to support tasks such as novelty detection in the application phase of the classifier. Other applications are possible, e.g., pruning at the end of a training process, extraction and ranking of rules contained in a classifier (see, e.g., [20]), concept drift or obsoleteness detection.

In our future work, we will also focus on measures for a quantification of experience gained by applying the knowledge extracted from data (e.g., by assessing its usefulness as in [19, 18]), extending the VI training techniques to apply them to large data sets, and on applications of the measures, e.g., in the field of collaborative intrusion detection in cyber-physical systems [29].

Further, we will investigate how our measures can be adjusted to suite other types of distributions (i.e., distributions that are required for other attribute types) as well.

Acknowledgment

This work was supported by the German Research Foundation (DFG) under grant SI 674/9-1 (project CYPHOC).

References

[1] Fubini theorem. Online, `http://www.encyclopediaofmath.org/index.php?title=Fubini_theorem&oldid=33052`, (last access: 13.05.2016)

[2] Adomavicius, G., Tuzhilin, A.: Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. IEEE Transactions on Knowledge and Data Engineering 17(6), 734–749 (2005)

[3] Asuncion, A., Newman, D.: UCI Machine Learning Repository. `http://archive.ics.uci.edu/ml/`, [Online] (last access 05/10/2016)

[4] Atzmueller, M., Baumeister, J., Puppe, F.: Rough-fuzzy MLP: modular evolution, rule generation, and evaluation. In: Proceedings of the 15th International Conference of Declarative Programming and Knowledge Management (INAP-2004). pp. 203–213. Potsdam, Germany (2004)

[5] Basu, S., Mooney, R.J., Pasupuleti, K.V., Ghosh, J.: Evaluating the novelty of text-mined rules using lexical knowledge. In: Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2001). pp. 233–238. San Francisco, CA (2001)

[6] Bishop, C.M.: Pattern Recognition and Machine Learning. Springer, New York, NY (2006)

[7] Bishop, C.: Novelty detection and neural network validation. IEE Proceedings – Vision, Image, and Signal Processing 141(4), 217–222 (1994)

[8] Bourassa, M., Fugère, J., Skillicorn, D.: Interestingness – directing analyst focus to significant data. In: Proceedings of the 2011 European Intelligence and Security Informatics Conference (EISIC 2011). pp. 200–207. Athens, Greece (2011)

[9] Cao, L.: Domain-driven data mining: Challenges and prospects. IEEE Transactions on Knowledge and Data Engineering 22(6), 755–769 (2010)

[10] Cao, L., Zhao, Y., Zhang, H., Luo, D., Zhang, C., Park, E.: Flexible frameworks for actionable knowledge discovery. IEEE Transactions on Knowledge and Data Engineering 22(9), 1299–1312 (2010)

[11] Chang, C., Lin, C.: LIBSVM: a library for support vector machines. ACM transactions on intelligent systems and technology (2012)

[12] Cover, T.M., Hart, P.E.: Nearest neighbor pattern classification. IEEE Transactions on Information Theory 13(1), 21–27 (1967)

[13] Di Fiore, F.: Visualizing interestingness. In: Zanasi, A., Brebbia, C., Ebecken, N., Melli, P. (eds.) Data Mining III. WIT Press, Southampton, U.K. (2002)

[14] Dubois, D., Hullermeier, E., Prade, H.: A systematic approach to the assessment of fuzzy association rules. Data Mining and Knowledge Discovery 13(2), 167–192 (2006)

[15] Duda, R.O., Hart, P.E., Stork, D.G.: Pattern Classification. John Wiley & Sons, Chichester, New York, NY (2001)

[16] Embrechts, M.J., Szymanski, B., Sternickel, K.: Introduction to scientific data mining: Direct kernel methods and applications. In: Ovaska, S.J. (ed.) Computationally Intelligent Hybrid Systems: The Fusion of Soft Computing and Hard Computing, chap. 10, pp. 317–362. IEEE Press, Hoboken, NJ (2005)

[17] Fayyad, U.M., Piatetsky-Shapiro, G., Smyth, P.: Knowledge discovery and data mining: Towards a unifying framework. In: Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining (KDD 1996). pp. 82–88. Portland, OR (1996)

[18] Fisch, D., Jänicke, M., Kalkowski, E., Sick, B.: Techniques for Knowledge Acquisition in Dynamically Changing Environments. ACM Transactions on Autonomous and Adaptive Systems 7(1) (2011)

[19] Fisch, D., Jänicke, M., Kalkowski, E., Sick, B.: Learning from others: Exchange of classification rules in intelligent distributed systems. Artificial Intelligence 187–188, 90–114 (2012)

[20] Fisch, D., Kalkowski, E., Sick, B., Ovaska, S.J.: In your interest: Objective interestingness measures for a generative classifier. In: Proceedings of the 3rd International Conference on Agents and Artificial Intelligence (ICAART 2011). pp. 414–423. Rome, Italy (2011)

[21] Fisch, D., Kühbeck, B., Sick, B., Ovaska, S.J.: So near and yet so far: New insight into properties of some well-known classifier paradigms. Information Sciences 180(18), 3381–3401 (2010)

[22] Fisch, D., Sick, B.: Training of radial basis function classifiers with resilient propagation and variational Bayesian inference. In: Proceedings of the International Joint Conference on Neural Networks (IJCNN '09). pp. 838–847. Atlanta, GA (2009)

[23] Fisch, D.: Intelligente technische Systeme mit der Fähigkeit zum kollaborativen Wissenserwerb. No. 1 in Intelligent Embedded Systems, kassel university press, Kassel, Germany (2012)

[24] Fisch, D., Kalkowski, E., Sick, B.: Knowledge fusion for probabilistic generative classifiers with data mining applications. Knowledge and Data Engineering, IEEE Transactions on 26(3), 652–666 (2014)

[25] Fubini, G.: Sugli integrali multipli. Rend. Acc. Naz. Lincei 16, 608–614 (1907)

[26] Glass, D.H.: Entailment and symmetry in confirmation measures of interestingness. Information Sciences 279, 552–559 (2014)

[27] Greco, S., Słowiński, R., Szczęch, I.: Properties of rule interestingness measures and alternative approaches to normalization of measures. Information Sciences 216, 552–559 (2012)

[28] Gruhl, C.: Self-adapting generative modeling techniques – a basic building block for many organic computing techniques. In: Tomforde, S., Sick, B. (eds.) Organic Computing Doctoral Colloquium 2014, pp. 99–110. kassel university press, Kassel, Germany (2014)

[29] Hähner, J., Rudolph, S., Tomforde, S., Fisch, D., Sick, B., Kopal, N., Wacker, A.: A concept for securing cyber-physical systems with organic computing techniques. In: Proceedings of the 26th International Conference on Architecture of Computing Systems (ARCS 2013). pp. 1–12. Prague, Czech Republic (2013)

[30] Hebert, C., Cremilleux, B.: A unified view of objective interestingness measures. In: Perner, P. (ed.) Machine Learning and Data Mining in Pattern Recognition, pp. 533–547. No. 4571 in LNAI, Springer, Berlin, Heidelberg, Germany (2007)

[31] Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating collaborative filtering recommender systems. ACM Transactions on Information Systems 22(1), 5–53 (2004)

[32] Hilderman, R.J., Hamilton, H.J.: Evaluation of interestingness measures for ranking discovered knowledge. In: Proceedings of the 5th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD 2001). pp. 247–259. Kowloon, Hong Kong, China (2001)

[33] Hilderman, R.J., Hamilton, H.J.: Knowledge Discovery and Measures of Interest. Kluwer Academic Publishers, Norwell, MA (2001)

[34] Jalali-Heravi, M., Zaiane, O.: A study on interestingness measures for associative classifiers. In: Proceedings of the 2010 ACM Symposium on Applied Computing. pp. 1039–1046. Sierre, Switzerland (2010)

[35] Jaroszewicz, S., Simovici, D.A.: Interestingness of frequent itemsets using Bayesian networks as background knowledge. In: Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2004). pp. 178–186. Seattle, WA (2004)

[36] Ji, S., Krishnapuram, B., Carin, L.: Variational Bayes for continuous hidden Markov models and its application to active learning. IEEE Transactions on Pattern Analysis and Machine Intelligence 28(4), 522–532 (4 2006)

[37] Jin, Y., Sendhoff, B.: Extracting interpretable fuzzy rules from rbf networks. Neural Processing Letters 17(2), 149–164 (2003)

[38] Kannan, S., Bhaskaran, R.: Association rule pruning based on interestingness measures with clustering. International Journal of Computer Science Issues 6(1), 35–43 (2009)

[39] Kretschmann, E., Fleischmann, W., Apweiler, R.: Automatic rule generation for protein annotation with the c4.5 data mining algorithm applied on swiss-prot. Bioinformatics 17(10), 920–926 (2001)

[40] Lenca, P., Meyer, P., Vaillant, B., Lallich, S.: On selecting interestingness measures for association rules: User oriented description and multiple criteria decision aid. Computing, Artificial Intelligence and Information Management 184, 610–626 (2008)

[41] Lewis, D.D., Gale, W.A.: A sequential algorithm for training text classifiers. In: Proceedings of the Seventeenth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '94). pp. 3–12. Dublin, Ireland (1994)

[42] Masood, A., Ouaguenouni, S.: Probabilistic measures for interestingness of deviations – a survey. International Journal of Artificial Intelligence & Applications (IJAIA) 4(2), 39–61 (2013)

[43] Masood, A., Soong, S.: Measuring interestingness – perspectives on anomaly detection. Computer Engineering and Intelligent Systems 4(1), 29–39 (2013)

[44] McGarry, K.: A survey of interestingness measures for knowledge discovery. The Knowledge Engineering Review 20(1), 39–61 (2005)

[45] McLachlan, G., Peel, D.: Finite Mixture Models. John Wiley & Sons, New York, NY (2000)

[46] MNIST: MNIST handwritten digit database. http://yann.lecun.com/exdb/mnist/ (last access 05/18/2016)

[47] Padmanabhan, B., Tuzhilin, A.: Unexpectedness as a measure of interestingness in knowledge discovery. Decision Support Systems 27(3), 303–318 (1999)

[48] Padmanabhan, B., Tuzhilin, A.: On characterization and discovery of minimal unexpected patterns in rule discovery. IEEE Transactions on Knowledge and Data Engineering 18(2), 202–216 (2006)

[49] Piatetsky-Shapiro, G., Matheus, C.: The interestingness of deviations. In: Proceedings of the AAAI-94 Workshop on Knowledge Discovery in Databases (KDD 1994). pp. 25–36. Seattle, WA (1994)

[50] Platt, J.C.: Fast training of support vector machines using sequential minimal optimization. In: Advances in Kernel Methods - Support Vector Learning. MIT Press (January 1998), http://research.microsoft.com/apps/pubs/default.aspx?id=68391, [Online] (last access 05/10/2016)

[51] Quinlan, R.: C4.5: Programs for Machine Learning. Morgan Kaufmann Publishers, San Mateo, CA (1993)

[52] Reitmaier, T., Calma, A.: Resp-$k$nn: A semi-supervised classifier for sparsely labeled data in the field of organic computing. In: Organic Computing Doctoral Dissertation Colloquium. vol. 4, pp. 85–97. kassel university press GmbH (2014)

[53] Reitmaier, T., Sick, B.: Let us know your decision: Pool-based active training of a generative classifier with the selection strategy 4DS. Information Sciences 230, 106–131 (2013)

[54] Reitmaier, T., Calma, A., Sick, B.: Transductive active learning – a new semi-supervised learning approach based on iteratively refined generative models to capture structure in data. Information Sciences 293, 275–298 (2015)

[55] Reitmaier, T., Sick, B.: The responsibility weighted mahalanobis kernel for semi-supervised training of support vector machines for classification. Information Sciences 323, 179–198 (2015)

[56] Ripley, B.: Pattern recognition and neural networks. http://www.stats.ox.ac.uk/pub/PRNN/, [Online] (last access 05/10/2016)

[57] Shaharanee, I., Hadzic, F., Dillon, T.: Interestingness measures for association rules based on statistical validity. Knowledge-Based Systems 24, 386–392 (2011)

[58] Shekar, B., Natarajan, R.: A framework for evaluating knowledge-based interestingness of association rules. Fuzzy Optimization and Decision Making 3(2), 157–185 (2004)

[59] Suzuki, E.: In pursuit of interesting patterns with undirected discovery of exception rules. In: Arikawa, S., Shinohara, A. (eds.) Progress in Discovery Science, pp. 504–517. No. 2281 in LNAI, Springer, Berlin, Heidelberg, Germany (2002)

[60] Taha, I., Ghosh, J.: Evaluation and ordering of rules extracted from feedforward networks. In: Proceedings of the International Conference on Neural Networks. vol. 1, pp. 408–413. Houston, TX (1997)

[61] Tew, C., Giraud-Carrier, C., Tanner, K., Burton, S.: Behavior-based clustering and analysis of interestingness measures for association rule mining. Data Mining and Knowledge Discovery 27(2), 1004–1045 (2013)

[62] Tuzhilin, A.: Knowledge evaluation: Other evaluations: usefulness, novelty, and integration of interesting news measures. Handbook of Data Mining and Knowledge Discovery 1, 496–508 (2002)

[63] UCI: Kdd cup 1999 data. http://kdd.ics.uci.edu/databases/kddcup99/kddcup99.html, [Online] (last access 05/10/2016)

[64] UCL: UCL/MLG Elena Database. http://mlg.info.ucl.ac.be/index.php?page=Elena, [Online] (last access 09/18/2015)

[65] Vashishtha, J.: Interestingness measures in rule mining: A valuation. Garima Int. Journal of Engineering Research and Applications 4(7), 93–100 (2014)

[66] Weng, J., Lim, E.P., He, Q., Leung, C.K.: What do people want in microblogs? measuring interestingness of hashtags in twitter. In: Proceedings of the 10th International Conference on Data Mining (ICDM '10). pp. 1121–1126. Sydney, Australia (2010)